VIRTUAL EXPERIENCE
OCTOBER 11-14

2021 Fall Technical Forum
SCTE • NCTA • CABLELABS

Node Health Within Cox ACOE's Service Health Framework

Improve the health and quality of HFC services through predictive analytics

A Technical Paper prepared for SCTE by

Shane Yates
Executive Director
Cox Communications
6305 Peachtree Dunwoody rd.
Atlanta, GA 30328
888-566-7751
Shane.Yates@Cox.com

Brian Stublen
Director
Cox Communications
6305 Peachtree Dunwoody rd.
Atlanta, GA 30328
888-566-7751
Brian.Stublen@Cox.com

Alexis Hwang
Manager
Cox Communications
6305 Peachtree Dunwoody rd.
Atlanta, GA 30328
888-566-7751
Alexis.Hwang@Cox.com

Table of Contents

Title **Page Number**

1. Introduction ... 3
2. History ... 3
3. Service Health Framework ... 5
4. Node Health ... 5
5. How Does Node Health RF Work? ... 7
   - 5.1. Overview ... 7
   - 5.2. BAU Process (Previous Approach) ... 7
   - 5.3. Problem Statement ... 8
   - 5.4. Solution Proposal ... 10
   - 5.5. Conclusion and Business Impact ... 11
6. Putting the Model into Action: AWS Cloud Architecture ... 11
7. Further Discussion ... 12

Abbreviations ... 13

List of Figures

Title **Page Number**

Figure 1 – OSP7 Measurement System ... 3
Figure 2 – Three Pillars of Node Health ... 6
Figure 3 – Node Health Projects ... 6
Figure 4 – Node Health Models ... 7
Figure 5 – Decomposition of the Impact of a Degradation RF performance Node ... 8
Figure 6 – Correlation Study betweeen Proactive Maintenance and Node Health ... 9
Figure 7 – Correlation with Node Health and Transactions ... 9
Figure 8 – Relationship between Node Health and market-level transaction level/variation ... 10
Figure 9 – Model of improved customer interactions post-Node Health models ... 11
Figure 10 – Cloud-based Node Health architecture ... 12

List of Tables

Title **Page Number**

Table 1 – OSP7 Program Metrics ... 4

1. Introduction

The Service Health framework is a suite of analytic models designed to predict and prevent customer impacting issues, while improving overall service quality. While this framework includes a multitude of Network, Customer, Premise, etc. focused components, this project and abstract are focused on Node Health and improving the effectiveness of the workforce that operates the HFC network, the Field Network Technician (FNT).

The FNT group is responsible for responding to all network impacting events, restoring customers'services, and performing proactive network maintenance to prevent issues before they occur. Previously, this work was created through a combination of rules-based telemetry processing and manually generated events. While this approach provided insight into clear rules violations regarding customer service availability, performance, and utilization, it was both limited in its ability to find novel patterns impacting customers and in its effectiveness with prioritizing work based on the potential impact to customers.

However, through the Node Health Framework, Cox can analyze thousands of combinations and permutations of RF patterns that may impact customer's services now and into the future. Additionally, this framework uses the predicted impact to prioritize work, both daily at a macro-level and hourly at a micro-level. As a result, FNTs have increased the volume of proactive network maintenance, prevented thousands of customer impacting events, reduced transactions to record low levels (note: Cox has attributed a 10%-15% reduction in transactions from the Node Health initiatives), and radically improved customer service quality.

2. History

Network Health is an initiative with roots going back to 2015 within Cox Communications. This started as a KPI-driven initiative to measure the health of the HFC network. The ultimate output of the early initiative was a measurement system called OSP7 (pictured below).

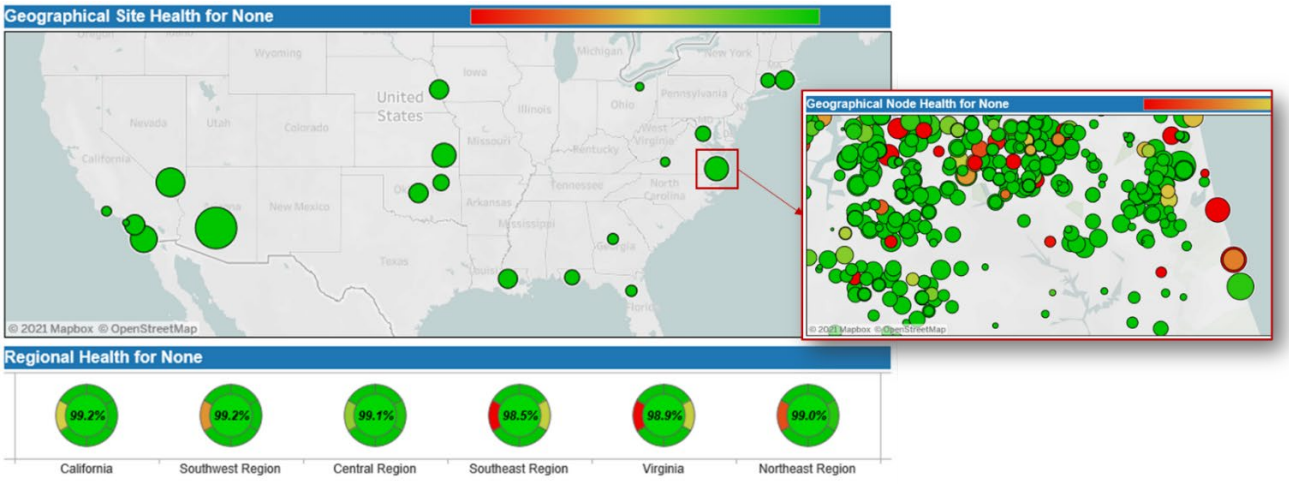

Figure 1 – OSP7 Measurement System

This measurement system began as a preliminary set of KPIs that measured a wide array of metrics, 35 to be specific, which ranged from operational performance to technician throughput/effectiveness. However, the analytics team believed there was room for improvement, so they put these KPIs through the two-factor test below:

1. Does this metric correlate with operational processes and practices? Or to put this more simply, can the Field Network Technician (FNT) organization successfully influence the metric either positively or negatively?
2. Does the metric correlate with business and financial objectives?

As a result of this two-factor test they created the OSP7 (or Outside Plant 7), which contained the 7 metrics represented below that correlated best with both FNT labor processes and business/financial objectives. The table below represents the first draft of the OSP7 program:

Table 1 – OSP7 Program Metrics

	Metric	Description	Spec
1	**DS RX (Downstream Receive)**	Downstream Power-level at the Customer modem within Docsis Spec	-12dB to +15
2	**US TX (Upstream Transmit)**	Upstream Power-level at the Customer modem within Docsis Spec	+30dB to +52
3	**DS SNR (Downstream Signal to Noise)**	Downstream SNR ration at the Customer modem within Docsis Spec	>+32 dB
4	**US SNR (Upstream Signal to Noise)**	Upstream SNR ration at the Customer modem within Docsis Spec	>+30 dB
5	**DS FEC (Downstream Forward Error Correction)**	Downstream packet loss at the Customer modem	<1% packet loss
6	**US FEC (Upstream Forward Error Correction)**	Upstream packet loss at the Customer modem	<1% packet loss
7	**Aggregate Health**	The combined score between all metrics	n/a

Within this program, the magic was not derived from the KPIs themselves, but rather from the user's ability to navigate the map (pictured above) and quickly identify problematic areas that required more focus from the FNT team. The FNT was now able to make decisions based on analytical modeling (i.e. what is healthy in green and unhealthy in red) provided to them using simple visualization principles. As a result, Cox experienced a substantial improvement in overall network performance (~1% improvement to OSP7 over 2 years associated with this work) and a corresponding drop in transactions (~5% reduction to Tech Support calls and UHT truck rolls). This reduction was calculated using several controlled experiments, where we measured improvements at the node, headend, and market/site level vs. control groups to validate and quantify the effect with precision.

3. Service Health Framework

As the organization matured, there was a clear need to evolve the space of analytics. As a result, Cox announced the formation of the Analytics Center of Excellence (ACOE) in 2017, which centralized all analytics teams across Cox's Technology and Operations organizations. With the formation of this new team, the ACOE launched new advanced analytical capabilities focused on the health and quality of the customers' services. This led to the creation of the Service Health Framework.

The Service Health Framework is based on three principles:

1. Most customer-impacting events can be predicted
2. Customer-impacting events are complex but unique and identifiable
3. We can shift from the customer as a diagnostic to prediction as a diagnostic

Additionally, there are four key focus areas that shape the framework:

1. Deep understanding of data that describes quality and use of service by our customers
2. Ability to predict future service impacts and mitigate or prevent them
3. Machines executing the right actions to take, at scale
4. Continuous learning capabilities to improve over time

From the intersection of these principles and focus areas, the ACOE designed the following structure to define the Service Health framework:

- **Node Health**: Predict, isolate, and prevent any customer experience issues at the node level
- **Location Health**: Predict home/business wiring, device, Wi-Fi customer experience issues
- **Usage Health**: Understand customer experience impact based on individual customer usage and behavior
- **Interaction Health**: Link human conversation to Service Health root cause

4. Node Health

As the Service Health framework launched, the previous foundations of network health, outlined in the history above, were used to quickly advance the overall framework. In particular, the lessons learned from using the OSP7 methodology/dashboard on how to direct labor, became an essential component for how Cox's ACOE developed the Node Health component of the Service Health framework. To help frame this up, the ACOE developed three pillars as the core to the Node Health concept:

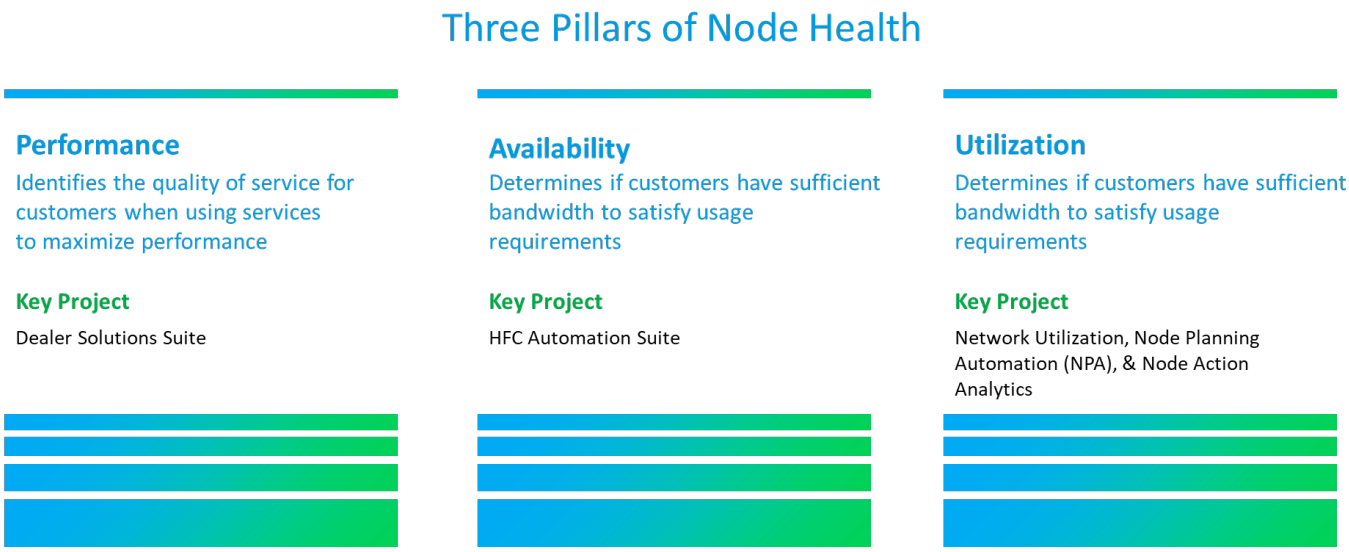

Figure 2 – Three Pillars of Node Health

These three pillars are the backbone of the Node Health framework and contain a suite of projects that seek to predict problems before customers recognize issues and prevent unnecessary transactions. Projects within these pillars span from predictive analytics, proactive customer resolution, and real-time root cause analysis. To dig a little deeper, there are six overarching projects that will deliver the core of the value within the Node Health framework.

These projects are described in the diagram below:

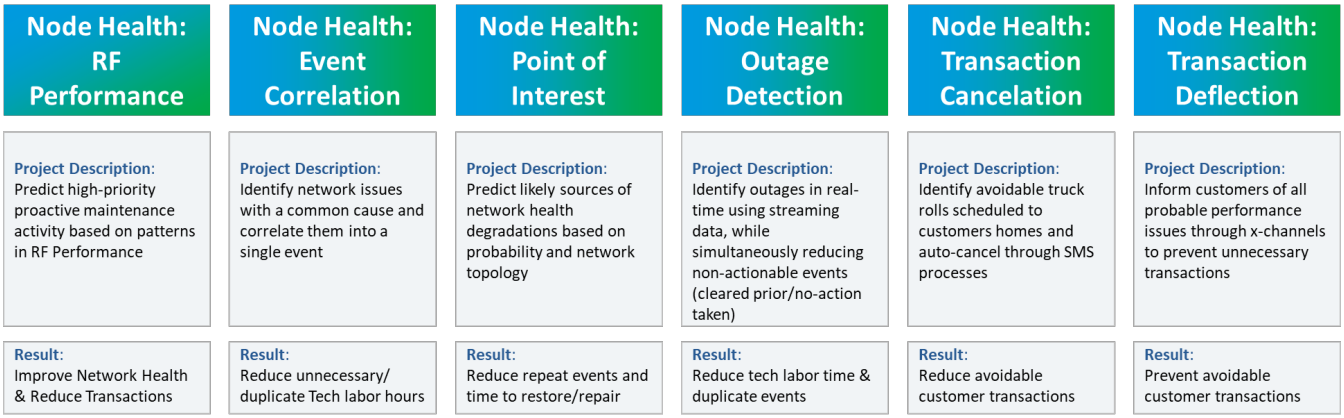

Figure 3 – Node Health Projects

The first program, titled "Node Health: RF Performance", is a proactive network maintenance model that takes in data from a multitude of sources and makes a prediction about the most impactful work that a technician can perform. Specifically, the model makes a prediction about the impact/value of performing work in every HFC node by creating patters based on historical node performance and the response/improvement those patterns recognized following proactive maintenance. Following this

prediction, the model then stack ranks those nodes based on the calculated value and routes the work according to priority to ensure the most impactful work is driven to the top priority. To make these predictions, we use the following suite of models listed below. One component left out of the diagram below is the continuous learning component of the model, which continues to learn how RF impairment patterns are impacted by proactive network maintenance activity to retrain and optimize the model.:

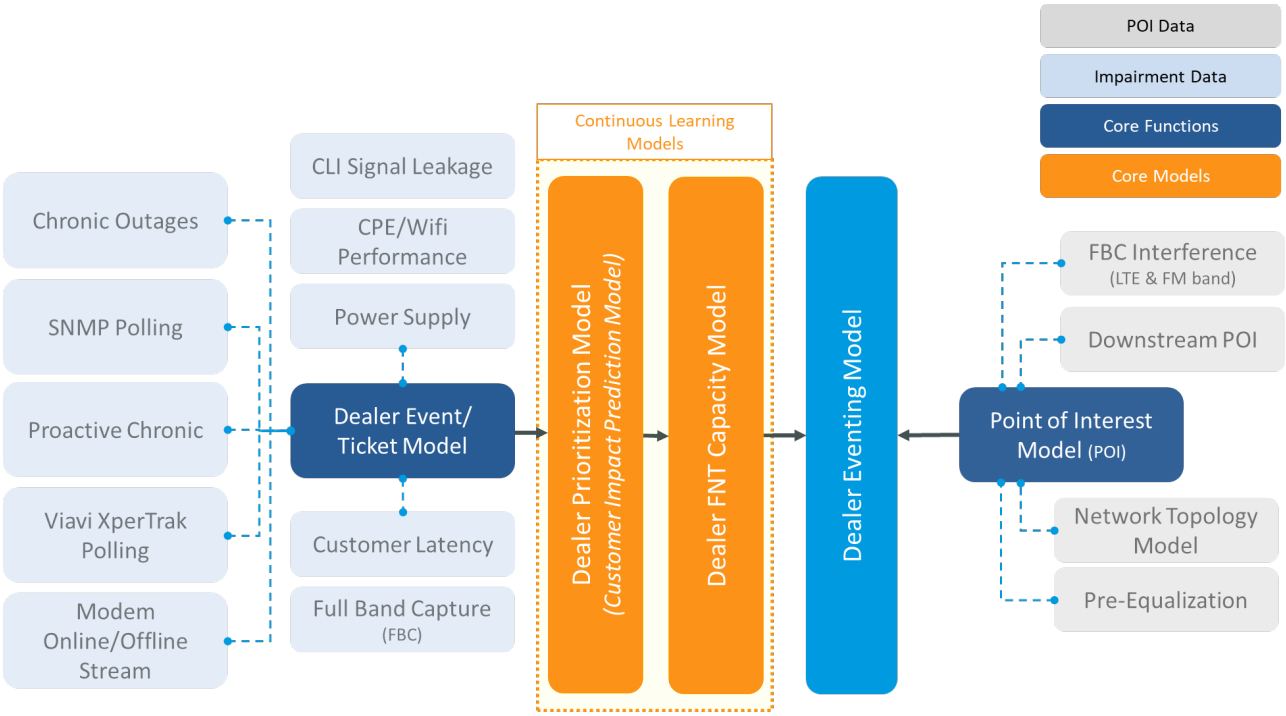

Figure 4 – Node Health Models

5. How Does Node Health RF Work?

5.1. Overview

Node Health RF Performance is a proactive platform for identifying areas of the network with recurring, intermittent and service affecting issues in order to optimize the use of Field Network Technicians' (FNT) time and labor before customers recognize an issue. With the introduction of Node Health RF Performance, Cox moves away from leveraging our customers as a primary diagnostic to leveraging machine learning models and their predictions as a primary diagnostic.

5.2. BAU Process (Previous Approach)

Before Node Health, our journey started with a customer calling in an issue. The customer in this case is uninformed throughout the entire process. Painfully, our technicians and care agents also have no visibility into the status of the issue. As a result of this reactive process, most operators have a difficult time enabling technicians or leveraging enhanced tooling to monitor and respond to the network events. Furthermore, managing customer interactions during the process is quite challenging as none of the various channels (web, app, phone, etc.) have visibility to these impairments. Below is a diagram outlining a recent example:

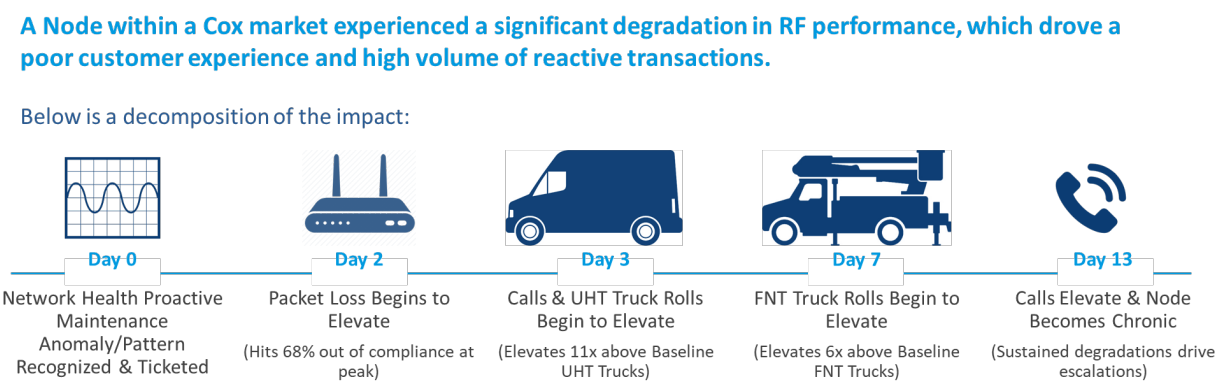

Figure 5 – Decomposition of the Impact of a Degradation RF performance Node

5.3. Problem Statement

Data shows a strong correlation between the amount of labor devoted to proactive maintenance event and the overall improvement in the health of the network (see diagram below, which depicts the correlation between proactive maintenance and network health). When network health is low, we see more outages, which leads to additional reactive work. However, the key challenge is how to prioritize this maintenance activity in a way that it can prevent the most transactions and can best improve customer experience. There are more proactive impairments to troubleshoot than there are technician labor hours in each week/month. As a result, much of the existing reactive work, both multi-customer outages and single-customer escalations, could be prevented if Cox successfully predicts which ***X*** impairments were the most critical and demanded work within ***Y*** days/hours.

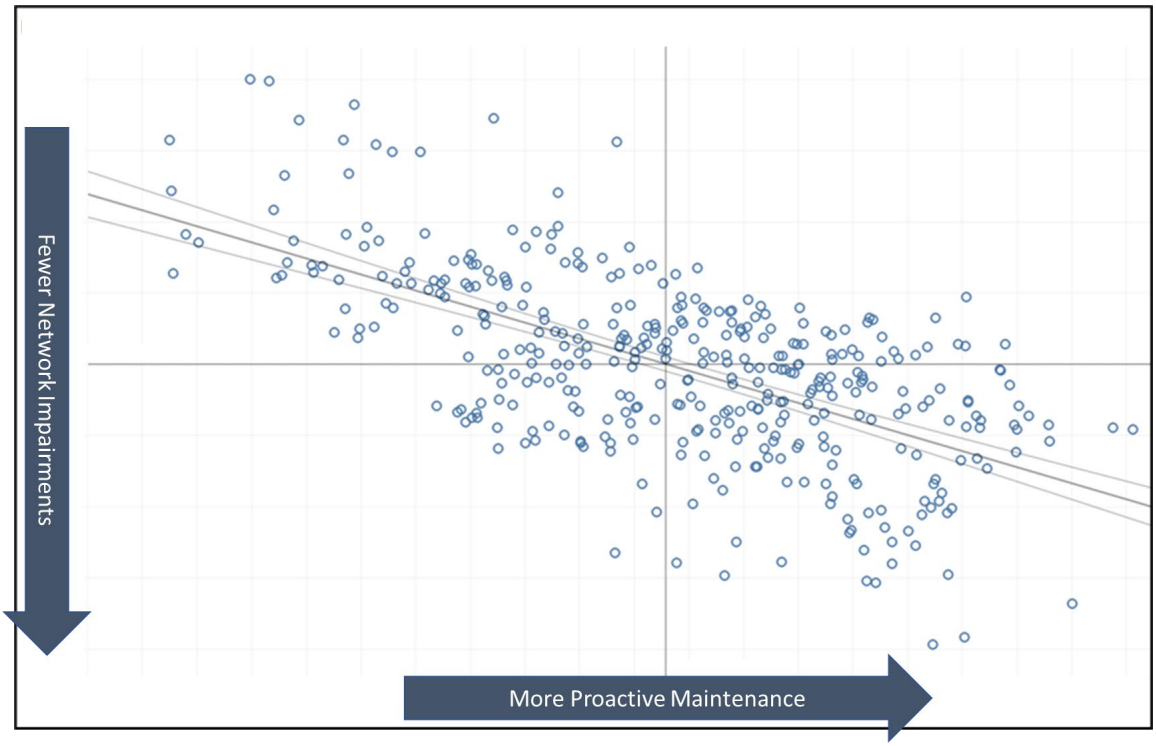

Figure 6 – Correlation Study betweeen Proactive Maintenance and Node Health

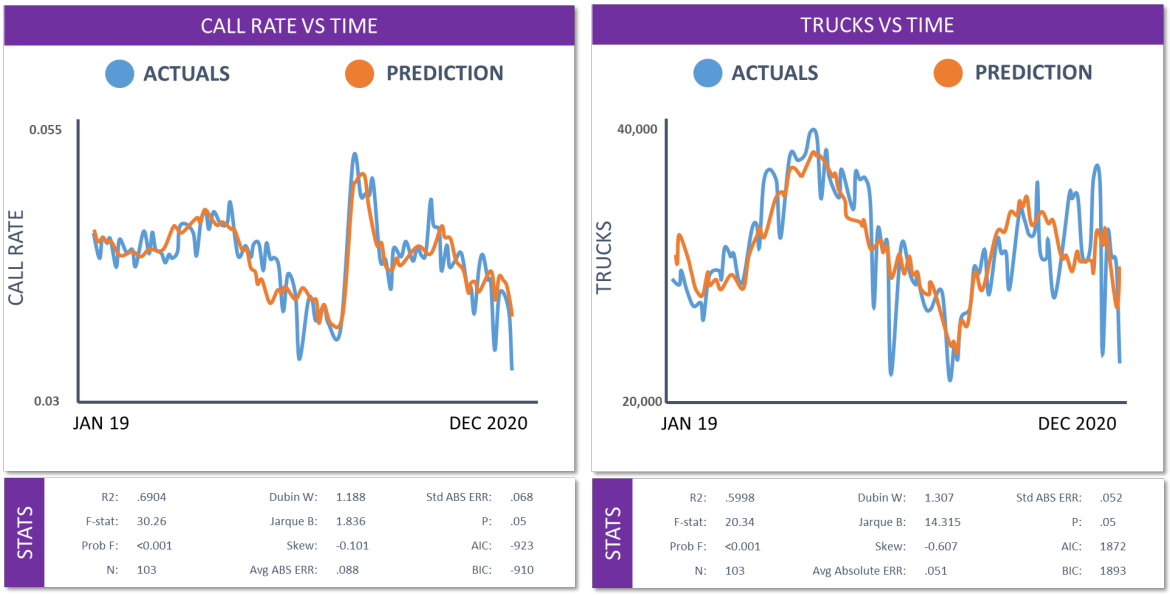

Figure 7 – Correlation with Node Health and Transactions

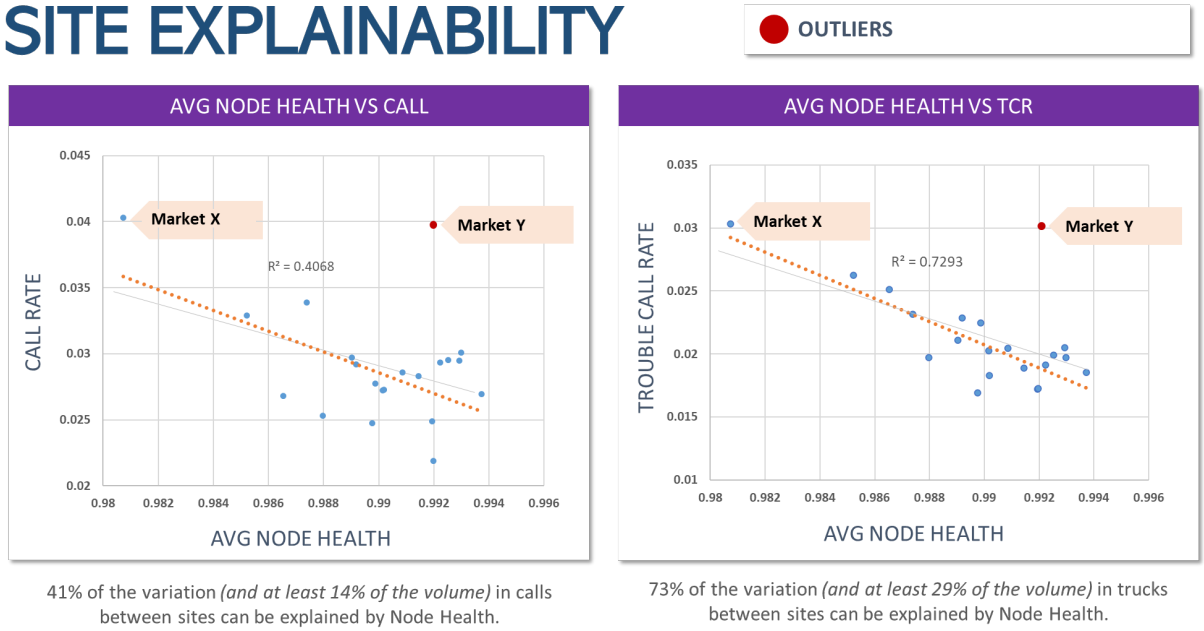

Figure 8 – Relationship between Node Health and market-level transaction level/variation

5.4. Solution Proposal

Cox uses a multitude of network and customer premise equipment (CPE) telemetry when identifying an issue. Specifically, they use the combination of sources highlighted in the diagram above, which retrieves the RF (Radio Frequency Signal data) and impairment metrics for multiple customer and network devices. By leveraging RF signal data and impairment data, devices with statistically similar patterns can be grouped together to act as an indicator of an issue on the network. Also, by leveraging real-time signal level and the frequency of occurrence of rapidly changing patterns, we can predict a time of day when the issue is present, which allows a technician to isolate the issue and find regular/irregular patterns in signal.

As a part of the capacity model, Cox also predicts the daily volume of reactive tickets to identify the volume of available labor hours/technicians grouped by their skills and schedules. Specifically, this model allows the Field Network Technician (FNT) team to receive a volume of proactive network maintenance work that meet the available capacity.

Finally, Node Health RF Performance ensemble models proactively creates a prioritized list of events/tickets that identify nodes with recurring/intermittent issue or continuous issue. This list of events/tickets is then routed based on priority and available technician capacity, which is performed daily (note: this process can run every 15 minutes but is currently running daily in accordance with existing processes). As a result, we can recast the diagram above with the following remediation and response:

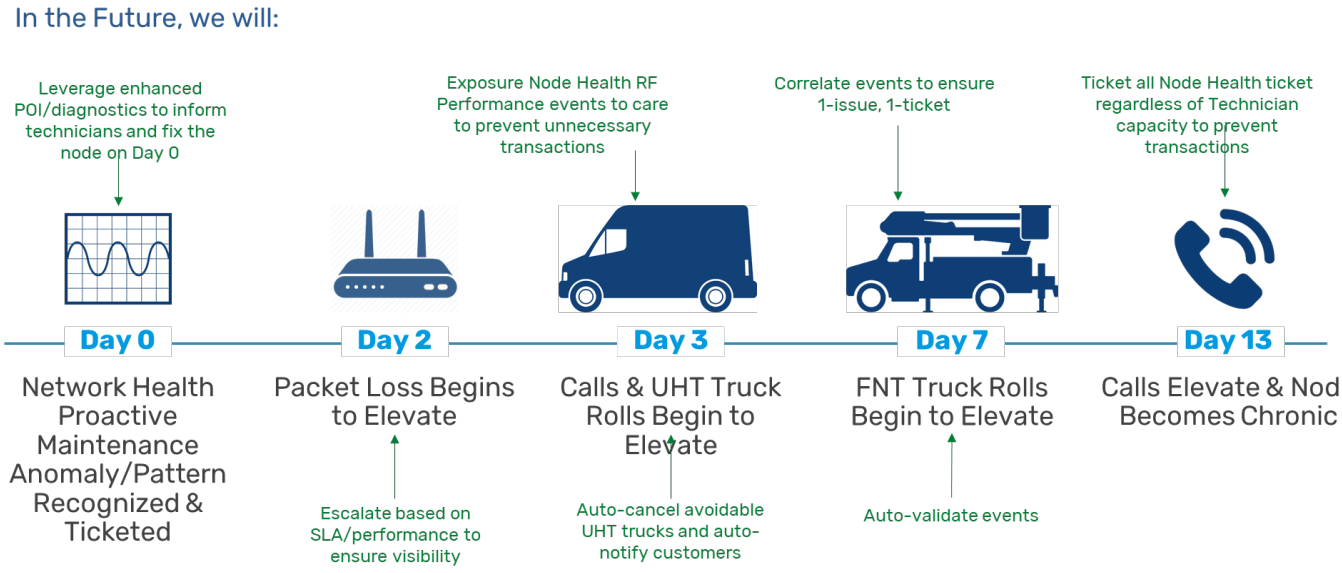

Figure 9 – Model of improved customer interactions post-Node Health models

5.5. Conclusion and Business Impact

The Node Health RF Performance application has been a significant driver in reducing the number of service-affecting events, multi-customer disruptions, and single-customer escalations. This has reduced unnecessary transactions, prevented customer service disruptions, and improved customer experience. Additionally, these capabilities have not only made technicians more efficient and effective, but they have also enabled next-generation communication with customers, allowing Cox to inform on potential issues and preemptively perform maintenance on parts of the network that are predicted to degrade. As a result of these ensemble models, Cox has experienced a ***5%*** reduction to Tech Support calls and in-home technician Service Call Truck Rolls.

6. Putting the Model into Action: AWS Cloud Architecture

The ACOE team used Amazon Web Services (AWS) as the primary cloud vendor for hosting this new technology. The following architecture diagram shows, generally, the services used by the Node Health RF Performance program.

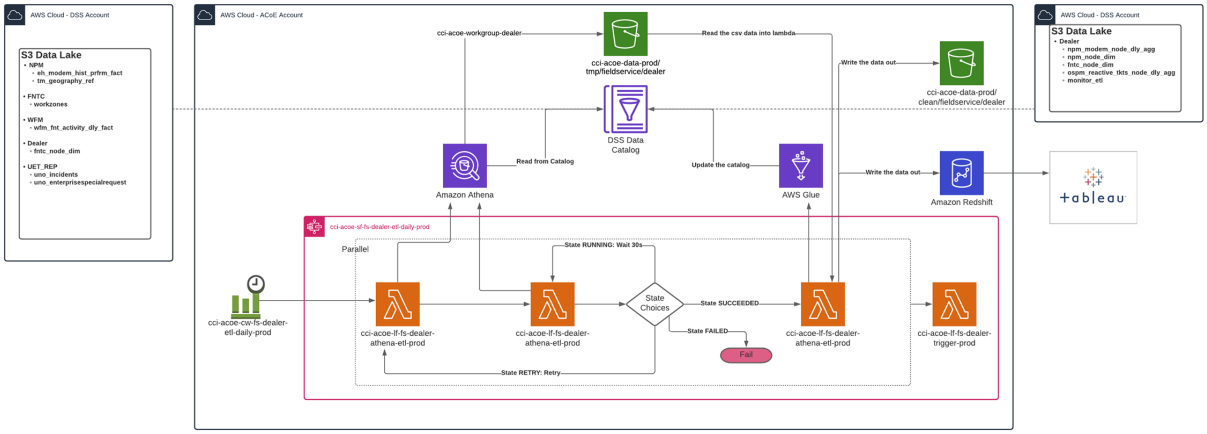

Figure 10 – Cloud-based Node Health architecture

When selecting to use a cloud partner, the ACOE had five primary considerations:

1. Scale: How do we scale our infrastructure easily when new projects/programs are approved?
2. Rapid Development: How do we get from ideation to production more quickly?
3. Next-Generation Capabilities: What platform gives Cox's ACOE capabilities that are multiple generations beyond existing tools?
4. Talent: What platform allows Cox's AOCE team to recruit and retain the best talent in the marketplace?
5. Future-Proof: What platform hardens our environment to ensure it is not obsolete in 2-5 years?

Moreover, the above architecture, which leverages Lamda, Athena, S3, Redshift, and Glue were selected for the following reasons:

- The overall cost of this pattern was the most cost effective model at the time
  - We don't need always-on database service since Node Health runs in batches periodically during the day rather than all of the time
  - The primary data sources were already in an S3 data lake
  - The tables are partitioned well such that we don't incur unnecessary cost when using the Athena service which charges by the GB scanned
- This pattern does not require ongoing maintenance of a database cluster
- The skills and experience required to create and manage Lambda functions are better aligned with the skills of our Data Scientists and Analysts over management of an EC2 instance.
- The Dealer application ETLs run (at the most frequent) once per hour and we did not want to incur the added time and money cost of turning up and down an EMR cluster (or EC2 Instance) or leaving one active
- Athena is a very fast service for its cost

7. Further Discussion

While this white paper focused on components of Node Health within the Cox ACOE Service Health framework, in future sessions we can focus on additional components of both Node Health and Service

Health. Areas of further discussion on Node Health that may be of interest and that are only briefly touched on in this whitepaper include:

- **Advanced event correlation**: How to improve grouping and correlation of related events based on predictive models and real-time root-cause analysis (RCA)
- **Point of interest (POI)**: How to predict potential drivers of problems within the node to reduce technician time to restore services and prevent repeat/unnecessary transactions
- **Outage detection**: How to reduce non-actionable events (i.e. often coded off as no-problem found/cleared prior), improve detection of true outages, and reduce the time to restore customers services using predictive analytics
- **Transaction cancelation**: How to cancel and prevent in-flight transactions using multiple models making predictions about customers' services and creating a dynamic interaction with the customer through SMS messaging
- **Transaction deflection**: How to improve predictions about service issues impacting customers and enrich the interaction through multiple channels to inform and educate the customer regarding restore/repair

These additional five projects augment the Node Health RF Performance project and add additional features that further improve the health of the network, reduce operating expense, and improve customer experience. Notably, these additional programs will reduce Tech Support calls and Service Call Truck Rolls (i.e. in-home technicians) by an incremental ***5%-10%***.

Abbreviations

RX	the receive power level to customers' modems
TX	the transmit power level from customers' modems
FEC	forward error correction
SNR	signal to noise ratio
US	upstream network performance
DS	downstream network performance
SCTE	Society of Cable Telecommunications Engineers
AWS	Amazon Web Services
POI	point of interest
ACOE	Analytics Center of Excellence
GB	gigabit
S3	AWS Simple Storage Services
EMR	AWS Elastic MapReduce
ETL	extract transform load (data transformation technique)
FNT	field network technician
CPE	customer premise equipment
RF	radio frequency
HFC	hybrid fiber-coax
OSP7	outside plant seven metrics system
UHT	universal home technician
KPI	Key Performance Indicator